Bayesian linking of geosynchronous orbital debris tracks as seen by the Large Synoptic Survey Telescope

M. D. Schneider

April 22, 2011

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Bayesian linking of geosynchronous orbital debris tracks as seen by the Large Synoptic Survey Telescope

Michael D. Schneider[a]

[a]Lawrence Livermore National Laboratory,
P.O. Box 808 L-210, Livermore, CA 94551-0808, USA

Abstract

We describe a Bayesian sampling model for linking and constraining orbit models from angular observations of "streaks" in optical telescope images. Our algorithm is particularly suited to situations where the observation times are small fractions of the orbital periods of the observed objects or when there is significant confusion of objects in the observation field. We use Markov Chain Monte Carlo to sample from the joint posterior distribution of the parameters of multiple orbit models (up to the number of observed tracks) and parameters describing which tracks are linked with which orbit models. Using this algorithm, we forecast the constraints on geosynchronous (GEO) debris orbits achievable with the planned Large Synoptic Survey Telescope (LSST). Because of the short 15 second exposure times, preliminary orbit determinations of GEO objects from LSST will have large and degenerate errors on the orbital elements. Combined with the expected crowded fields of GEO debris it will be challenging to reliably link orbital tracks in LSST observations given the currently planned observing cadence.

Keywords: space debris, orbit determination, large surveys

1. Introduction

The planned Large Synoptic Survey Telescope[1] (LSST) will image the entire visible sky every 3 nights for 10 years in a series of 15-second exposures covering a 9.6 square degree field of view to a limiting magnitude of 24.5 in the *r* band (LSST Science Collaborations and LSST Project, 2009). The combination of the regular and fast cadence with the wide field of view means the LSST may have the potential to map the distribution of orbital debris around the Earth to an unprecedented level. Earth-orbiting satellites with angular speeds less than 3.5 degrees per 15 seconds will appear as "streaks" in the individual exposures whose endpoints will yield angular positions of the satellites at the exposure start and end times. Objects in low-Earth orbit (LEO) will pass through the entire field of view in one exposure, but LSST could potentially measure the angular positions of tens of thousands of objects (with sizes down to a centimeter) near Geosynchronous altitudes (GEO). With the currently planned cadence, LSST will re-image some objects after about an hour, with the remaining objects being imaged again only after about 3 days. Here we investigate methods to constrain the orbital elements of GEO debris with the LSST.

Because the LSST goes so deep in each exposure, confusion of streaks across multiple exposures is a potential problem (see also DeMars & Jah, 2009). The streaks can be grouped according to length, reflecting angular speed, but there still could be several streaks in many exposures consistent with GEO objects if current GEO debris distribution forecasts are correct (e.g. Schildknecht et al., 2001, 2004). We can expect orbital debris to generally have albedo variegations and to be rotating in unknown ways. In principle, the temporal brightness fluctuations of distinct objects could be categorized and used as discriminating information (Payne et al., 2007). However, we expect that for many GEO objects the sampling rate of LSST observations will be too small to make the temporal brightness a useful quantity for linking uncorrelated tracks and we do not consider it further. Because 15 seconds of observation will yield very poor preliminary orbit determinations (PODs) for objects with periods around one day (Marsden, 1991; Gronchi, 2004), multiple streaks

---

*Corresponding author

Email address: schneider42@llnl.gov (Michael D. Schneider)

[1]http://www.lsst.org

in subsequent exposures could be erroneously linked yielding false entries in the debris catalogue. We therefore need a robust algorithm for POD uncertainty quantification to properly understand the probabilities for incorrect linking of tracks. It is also well known that linear methods (e.g. Kalman filter methods) for propagating orbit state vectors fail for observations with short arc lengths compared to the orbital period Bowell et al. (2002). So we also need a better method for propagating PODs to evaluate the linking probability between exposures. While the brute-force approach for linking tracks scales as the square of the number of uncorrelated tracks, the computational requirements for linking can be greatly mitigated by the use of KD-trees (Kubica et al., 2007; Granvik & Muinonen, 2008) or by the use of geometric methods to quickly evaluate the potential regions of overlap of propagated orbits (Milani et al., 2004). The main requirement on the linking algorithm is therefore robustness rather than computational speed.

In order to understand the capability of the LSST for constraining GEO debris orbits, we need algorithms for PODs and linking that can determine arbitrary error distributions for the PODs and propagate these errors in a robust way. One approach is to use parameterized error distributions that deviate from the Gaussian assumption (Muinonen & Bowell, 1993; DeMars et al., 2009). Alternatively, the requirement to obtain bound orbits given the angular positions and angular rates allows the specification of an "admissible region," even with poor data, that can then be compared for later track linking (Tommei et al., 2007; Maruskin et al., 2009; DeMars et al., 2010; Fujimoto & Scheeres, 2010). Under the "admissible region" formulation of the POD problem from optical observations, it is assumed that a series of angular coordinates of a satellite are available closely spaced in time from several back-to-back telescope exposures. This list of angular coordinates and times is combined to yield estimates of the angular position and angular rate of the satellite at a single epoch. The angular position and rate at a given epoch then provides bounds on the range and range rate with the requirement to obtain bound and stable orbits. In this paper we are considering the same problem that is addressed by the specification of the "admissible region" with two important differences. First, rather than combining the list of observed angular positions at closely spaced times into a single estimate of the angular position and rate, we use all the observed angular positions at the given times to jointly constrain the parameters of the orbit model for the observed object (as in Muinonen & Bowell, 1993). The orbit model parameters that are constrained can be chosen for convenience and could include, for example, the 3D state vector at a given epoch or the 6 Keplerian elements. Second, we link PODs from two separate observed tracks based on the overlap between the posterior probability distributions of the orbit model parameters given the track observations (Granvik, 2007) (which can be different from using the overlap in phase space as used when linking tracks based on the admissible regions).

To constrain the orbit model parameters observed angular positions over a short time interval we consider a common approach to characterizing arbitrarily complex error distributions via Markov Chain Monte Carlo (MCMC), which was recently applied to the problem of PODs with short arcs by Oszkiewicz et al. (2009). By following a random walk in the orbital elements MCMC allows us to generate samples from the joint posterior distribution of the orbital elements given a set of observed tracks with no constraints on the form of the posterior. (The constraint to obtain bound orbits can be imposed either by a restrictive choice of orbit model parameters, such as Keplerian elements, or by specifying prior distributions that down-weight parameter values that would lead to unbound orbits.) Because the MCMC sampling of the orbital elements is conditioned on a particular linking of the observed tracks, we can think of the linking problem as a model selection problem in this context. We are then able to draw on model selection frameworks in the literature to achieve the linking via MCMC. The practicality of this approach lies in finding efficient MCMC sampling algorithms. Our goals in this paper are then to establish the feasibility of linking uncorrelated tracks via Bayesian model selection algorithms and to apply this algorithm to forecast a "typical" linking probability for GEO debris with one night of observations with the LSST. We leave studies of the efficiency and scaling of the linking algorithm for later work.

This paper is organized as follows. We describe the MCMC methods for POD from short arc observations in Section 2.1. Our sampling model for simultaneous POD and linking of tracks is described in Section 2.2. We demonstrate the algorithms from Section 2 with an LSST case study in Section 3. We draw conclusions from the case study in Section 4. Throughout, we assume only the Newtonian gravitational force from the Earth and Keplerian orbits.

2. Method

In this section we describe a general algorithm for both POD and linking of uncorrelated tracks with no restrictions placed on the number of observed tracks or the times between their observations. In particular, we allow for arbitrary numbers of orbit wrappings, which is an important parameter to consider when evaluating all possible track linking probabilities within a dataset spanning many days or longer. In Section 3 we restrict the application of our algorithm to a single night of observations of GEO debris, but emphasize that the algorithm in this section is more generally applicable.

2.1. Bayesian orbit determination from optical observations

Individual exposures of the telescope will have satellite streaks whose endpoint(s), when visible in the exposure, give angular positions of the satellites at the exposure start

and/or end times. We therefore assume as input data a series of RA and DEC coordinates at times $\mathbf{t}$ along with associated angle uncertainties σ_{obs} for the streak endpoints extracted from a series of telescope exposures. When multiple telescope exposures are taken in a short time interval (e.g. within about a minute for GEO debris) we will assume the streaks in each exposure can be linked with perfect certainty and will refer to such a series of streaks as a "track" consisting of multiple angular measurements at known times. In general there is a degeneracy between the exposure start and end times unless the satellite orbital direction is known. Because we are considering LSST's sensitivity to geostationary orbits, which are typically prograde, we will ignore this degeneracy in what follows.

Assuming that the angular measurement uncertainties are uncorrelated between distinct tracks and that the uncertainties are Gaussian distributed gives the following likelihood for the data given an orbit model:

$$L_{1-\mathrm{orbit}}(\hat{\mathbf{y}}|\boldsymbol{\theta}) \propto$$
$$\exp\left[-\frac{1}{2}\sum_{i=1}^{N_{\mathrm{track}}}(\hat{\mathbf{y}}_i - \bar{\mathbf{y}}(\mathbf{t}_i, \boldsymbol{\theta}))^T \Sigma_i^{-1} (\hat{\mathbf{y}}_i - \bar{\mathbf{y}}(\mathbf{t}_i, \boldsymbol{\theta}))\right], \quad (1)$$

where N_{track} is the number of distinct observed tracks, $\mathbf{t}_i \equiv (t_{i,1}, t_{i,2})$ are the observation times of the streak endpoints in track i, $\hat{\mathbf{y}}_i \equiv (\mathtt{RA}(t_{i,1}), \mathtt{DEC}(t_{i,1}), \mathtt{RA}(t_{i,2}), \mathtt{DEC}(t_{i,2}))$ is the observed angular positions of the track at times $\mathbf{t}_i$, $\bar{\mathbf{y}}(t, \boldsymbol{\theta})$ is the model prediction for the satellite angular position at time t given orbital parameters $\boldsymbol{\theta}$, and Σ_i is the noise covariance matrix for the RA and DEC observations at times $\mathbf{t}_i$, which we will mostly assume to be $\Sigma_i = \sigma_{\mathrm{obs}}^2 \mathbf{I}_4$ (i.e. diagonal with the same errors for all t).

The posterior probability distribution for the orbital elements $\boldsymbol{\theta}$ given a series of observed tracks $\hat{\mathbf{y}}$ is derived from Eq. (1) via Bayes' theorem, $P(\boldsymbol{\theta}|\hat{\mathbf{y}}) \propto L_{1-\mathrm{orbit}}(\hat{\mathbf{y}}|\boldsymbol{\theta}) P(\boldsymbol{\theta})$, where $P(\boldsymbol{\theta})$ is the prior distribution for the orbital elements. We use Markov Chain Monte Carlo (MCMC) with Metropolis-Hastings (MH) updates to draw samples of $\boldsymbol{\theta}$ from $P(\boldsymbol{\theta}|\hat{\mathbf{y}})$. The MH algorithm requires the specification of a proposal distribution, $P_{\mathrm{MH}}(\boldsymbol{\theta})$. The MCMC chain is then advanced by first drawing a proposed step $\boldsymbol{\theta}'$ from $P_{\mathrm{MH}}(\boldsymbol{\theta})$ and then accepting the proposal with probability,

$$\alpha_{\mathrm{MH}} = \frac{P(\boldsymbol{\theta}'|\hat{\mathbf{y}}) P_{\mathrm{MH}}(\boldsymbol{\theta})}{P(\boldsymbol{\theta}|\hat{\mathbf{y}}) P_{\mathrm{MH}}(\boldsymbol{\theta}')}. \quad (2)$$

When the proposal is rejected, the chain remains at the current location $\boldsymbol{\theta}$.

The choice of $P_{\mathrm{MH}}(\boldsymbol{\theta})$ is critical to obtaining a large acceptance probability for each MH update and therefore an efficient sampling of the posterior for the orbital elements. When the errors in the PODs are large (as will be the case for a single LSST visit observing GEO objects), the orbital element posterior can be highly degenerate in any of the common orbital element parameterizations. Obtaining efficient MH updates then requires a proposal distribution that specifically accounts for these degenerate directions

in parameter space. Virtanen et al. (2001) have developed such a proposal distribution in a technique they call "statistical ranging." The idea behind statistical ranging is to parameterize the orbit using the two 3D position vectors (with origin at the center of the Earth) to the observed track endpoints. The angular components of these position vectors are tightly constrained by the observations and it is primarily the ranges that must be sampled (hence the name statistical ranging). We follow Virtanen et al. (2001) and specify independent Gaussian proposal distributions for each of the angular coordinates centered on the observed angular coordinates and a bivariate Gaussian distribution for the two ranges, with a large correlation coefficient (~ 0.99). The proposed angular coordinates will always be consistent with the observations, so a careful choice of the range proposal variance will lead to large MH acceptance probabilities no matter how degenerate the posterior may be in other parameterizations. Through trial and error we found that we could consistently achieve MH acceptance probabilities above 30% for modeled GEO debris when the range proposal standard deviation was set to,

$$\sigma_\rho \approx \frac{5000\,\mathrm{km}}{\Delta t/(34\,\mathrm{s})}, \quad (3)$$

where Δt is the time difference in seconds between the two observations defining the 3D position vectors. We also determined the correlation coefficient for the two range proposals according to,

$$\mathrm{cor}(\rho) \approx 1 - 2e\sin\left(\pi\frac{\Delta t}{1\,\mathrm{day}}\right), \quad (4)$$

where e is an initial guess for the orbit eccentricity (if the initial guess is not very good, the chain will still work but be somewhat less efficient). The range proposal mean can either be set to a fixed value at a typical range for the debris population being studied, or to the value of the range parameter at the current chain step. In our studies we found the former choice somewhat more efficient for starting the chains when the likelihood values are small (the so-called "burn-in" phase), while the latter setting for the range mean was marginally better for sampling near the peak of the likelihood. When the orbit model is being constrained with more than one streak, there is an ambiguity about which two streak endpoints to choose for parameterizing the orbit. Virtanen et al. (2001) advocate choosing the first and last observations ordered in time, but a random choice can work equally well in many cases.

While using two position vectors is a convenient way to parameterize $P_{\mathrm{MH}}(\boldsymbol{\theta})$, it will be desirable for other applications to parameterize the orbit using, e.g., keplerian or equinoctial elements (Broucke & Cefola, 1972). Shefer (2010) has recently derived the general conversion of two 3D position vectors to orbital elements for arbitrary conic sections and number of orbit wrappings. In general, the half plane of the difference of true anomalies at the times of the two position vectors, $\mathrm{sgn}(\cos(\nu_2 - \nu_1))$, and the number of orbit wrappings, $\lambda \in \mathbb{Z}_+$, must also be specified to

uniquely identify the two position vectors with a set of orbital elements. Because these quantities are not known from the 3D position vector proposals, we determine them iteratively by starting with guessed values and re-deriving the orbital elements from the position vectors until the orbital elements do not change. The guesses given two position vectors are,

$$n_{\text{orbits}} = (t_2 - t_1)/T_{\text{GEO}} \quad (5)$$

$$\lambda_{\text{guess}} = \text{floor}\,(n_{\text{orbits}}) \quad (6)$$

$$\text{sgn}(\cos(\nu_2 - \nu_1))_{\text{guess}} = \text{sgn}\,(n_{\text{orbits}} - \lambda_{\text{guess}} - 0.5)\,, \quad (7)$$

where $T_{\text{GEO}} = 86163$ seconds is a typical period of a GEO orbit. Once a preliminary set of orbital elements is found, updated values of λ and $\text{sgn}(\cos(\nu_2 - \nu_1))$ can be computed from the orbital period and true anomalies determined from the preliminary elements.

2.2. Linking tracks via Bayesian model selection

When tracks are observed in multiple exposures and the PODs from a single exposure cannot be propagated with high precision, then we need a sampling model that allows for the possibility that the tracks belong to distinct objects. So, in addition to constraining orbital elements, we also have a model selection and linking problem in determining how many unique objects are observed and which observations belong to which objects.

For N_{track} observed tracks, there could be between 1 and N_{track} unique objects generating the observations (in the absence of other identifying information such as the magnitudes of the streaks). We will assume sampling distributions for $N_{\text{orbit}} \leq N_{\text{track}}$ orbits such that the combined likelihood for the data given N_{orbit} orbit models is,

$$L\left(\hat{\mathbf{y}}\,|\,\{\boldsymbol{\theta}_j\}_{j=1}^{N_{\text{orbit}}}\right) = \prod_{j=1}^{N_{\text{orbit}}} \prod_{i=1}^{N_{\text{track}}} \left[L_{1-\text{orbit}}(\hat{\mathbf{y}}_i|\boldsymbol{\theta}_j)\right]^{k_j^i}\,, \quad (8)$$

where k_j^i is 1 if orbit j generated data point i and is zero otherwise. We call the k_j^i parameters "orbit selector" parameters. (See Hogg et al., 2010, for a simple example of similar model selection parameters.) With the combined likelihood of Eq. (8) we have reduced the model selection problem of linking orbits to a parameter estimation problem over the product space of all $(k_j^i, \boldsymbol{\theta}_j)$ (Carlin & Chib, 1995; Godsill, 2001). This is an extremely valuable simplification in that MCMC methods can be used to sample the parameter product space more efficiently than a brute force matching of all pairs of tracks because minimal computation will be wasted on sampling uncorrelated directions in the product space.

Because a given track can be associated with only one orbit model at each step in the MCMC chain the orbit selectors have the constraint,

$$\sum_{j=1}^{\min(i,N_{\text{orbit}})} k_j^i = 1 \,\forall\, i. \quad (9)$$

That is, for a given i the set $\left\{k_j^i\right\}_{j=1}^{N_{\text{orbit}}}$ are all zero except for one parameter equal to one; which selects the orbit model for observation i. Without loss of generality, we can set $k_1^1 = 1$ (implying $k_j^1 = 0$ for $j = 2, \ldots, N_{\text{orbit}}$). Additionally we can generally set $k_j^i = 0$ for $i < j$. This removes a degeneracy where we simply relabel the orbit model indices. So the set $\left\{k_j^i\right\}$ forms a lower-triangular matrix,

$$\mathsf{k} \equiv \begin{pmatrix} 1 & 0 & \ldots & 0 \\ k_1^2 & k_2^2 & \ldots & 0 \\ \vdots & \vdots & \ddots & 0 \\ k_1^{N_{\text{track}}} & k_2^{N_{\text{track}}} & \ldots & k_{N_{\text{orbit}}}^{N_{\text{track}}} \end{pmatrix}. \quad (10)$$

An appropriate prior distribution for the set of N_{orbit} parameters selecting one of N_{track} possible outcomes of a single trial is the multinomial distribution,

$$P(\mathsf{k}|P_{\text{orbit}}) = \prod_{j=1}^{N_{\text{orbit}}} \left(P_{\text{orbit}}^{ij}\right)^{k_j^i} \,\forall\, i = j, \ldots, N_{\text{track}}, \quad (11)$$

where P_{orbit}^{ij} is the prior probability that observation i is generated by orbit model j. Because we do not, in general, know these prior probabilities, we can marginalize over the P_{orbit}^{ij} parameters with a conjugate Dirichlet prior,

$$P(P_{\text{orbit}}|\alpha_{\text{orbit}}) = \prod_{j=1}^{N_{\text{orbit}}} \left(P_{\text{orbit}}^{ij}\right)^{\alpha_{\text{orbit}}^{ij}} \,\forall\, i = j, \ldots, N_{\text{track}}, \quad (12)$$

with parameters $\alpha_{\text{orbit}}^{ij}$ that can be individually tuned to be more or less informative for each data point.

The full set of parameters in our sampling model now includes the orbital elements for each of the N_{orbit} orbit models, the orbit selectors for deciding which tracks are linked with which orbit models, and the prior probabilities for each orbit selector,

$$\boldsymbol{\Phi} \equiv \left(\boldsymbol{\theta}_j, k_j^i, P_{\text{orbit}}^{ij}\right), \quad (13)$$

for $i = j, \ldots, N_{\text{track}}$ and $j = 1, \ldots, N_{\text{orbit}}$.

The conjugate priors for the orbit selectors and P_{orbit} suggest that we could update these parameters using Gibbs sampling in our MCMC chain for sampling $\boldsymbol{\Phi}$. This doesn't quite work for reasons we will explain in the next section, but it is instructive to consider how this would be implemented. The updates of $\boldsymbol{\Phi}$ would proceed as,

1. For each i, draw new orbit selectors, $\left\{k_j^i\right\}$, as a Gibbs update from the conditional distribution,

$$P(k^i|\hat{\mathbf{y}}, \boldsymbol{\theta}, P_{\text{orbit}}) = \prod_j \left(L_{1-\text{orbit}}(\hat{\mathbf{y}}_i|\boldsymbol{\theta}_j)\right)^{k_j^i} \cdot P(k^i|P_{\text{orbit}}^i). \quad (14)$$

This is a Multinomial distribution in k_j^i with probabilities $L_{1-\text{orbit}}(\hat{\mathbf{y}}_i|\boldsymbol{\theta}_j) \cdot P_{\text{orbit}}^{ij}$.

2. For each i draw new prior probabilities, P_{orbit}^i, as a Gibbs update from the conditional distribution,

$$P(P_{\mathrm{orbit}}^i|\hat{\mathbf{y}}, \boldsymbol{\theta}, k^i) = P(k^i|P_{\mathrm{orbit}}^i) \cdot P(P_{\mathrm{orbit}}^i|\alpha_{\mathrm{orbit}}^i). \tag{15}$$

This is a Dirichlet distribution in P_{orbit}^{ij} with parameters $k_j^i + \alpha_{\mathrm{orbit}}^{ij}$.

3. For each j update $\boldsymbol{\theta}_j$ with a MH update as described in Eq. (2).

The Gibbs updates of k and P_{orbit} are fast to evaluate. So the success of the algorithm relies on obtaining efficient mixing for step 3 so that the orbit selector updates in step 1 will step often leading to thorough exploration of the space of viable orbit models.

2.2.1. Orbit model priors

The sampling model for $\boldsymbol{\Phi}$ is not complete until we specify prior distributions for the orbital elements, $\boldsymbol{\theta}_j$. The priors must be chosen carefully because when $k_j^i = 0 \,\forall i$ the likelihood in Eq. (8) is independent of $\boldsymbol{\theta}_j$ and the prior on $\boldsymbol{\theta}_j$ will completely specify the conditional posterior. Following Godsill (2001) we use the composite prior,

$$P(\boldsymbol{\Theta}|k) = P(\boldsymbol{\Theta}_k|k)\,P(\boldsymbol{\Theta}_{-k}|\boldsymbol{\Theta}_k, k), \tag{16}$$

where $\boldsymbol{\Theta}_k \equiv \left\{\boldsymbol{\theta}_j | \sum_{i=j}^{N_{\mathrm{track}}} k_j^i > 0\right\}$ and $\boldsymbol{\Theta}_{-k}$ is the complementary set. The second term in Eq. (16) is called a "pseudo-prior" by Carlin & Chib (1995) because it can be chosen arbitrarily without affecting the marginal distributions for the remaining parameters due to the indepence of Eq. (8) to $\boldsymbol{\Theta}_{-k}$.

The pseudo-prior must be chosen carefully however to obtain efficient stepping in the MCMC chain in the orbit selector parameter directions. To see this consider the orbit selector Gibbs update in Eq. (14). The only way that a given k_j^i can have a significant probability to step from a value of zero to one is if $L_{1-\mathrm{orbit}}(\hat{\mathbf{y}}_i|\boldsymbol{\theta}_i) \lesssim 1$. But, if at the current step $k_j^i = 0 \,\forall i$, then $\boldsymbol{\theta}_j$ will have been sampled from the pseudo-prior and will most often give $L_{1-\mathrm{orbit}}(\hat{\mathbf{y}}_i|\boldsymbol{\theta}_j) \ll 1$ unless the pseudo-prior has been carefully selected.

Looking again at Eq. (14) we can see that the prior choice for $\boldsymbol{\theta}_j$ that will lead to the most efficient sampling in the orbit selectors is $P(\boldsymbol{\Theta}_{-k}|k) = L_{1-\mathrm{orbit}}(\hat{\mathbf{y}}_i|\boldsymbol{\theta}_j)$ for given i. This is a strange sort of prior because it depends on the observations, $\hat{\mathbf{y}}$. But as mentioned above, $P(\boldsymbol{\Theta}_{-k}|k)$ can be chosen arbitrarily. Because $\boldsymbol{\theta}_j$ can be fit to several observations (indexed by different i values), it does not make sense to specify the pseudo-prior for a given i. Instead we choose,

$$P(\boldsymbol{\Theta}_{-k}|k, \hat{\mathbf{y}}) = \prod_{j \in -k} L_{1-\mathrm{orbit}}(\hat{\mathbf{y}}_j|\boldsymbol{\theta}_j). \tag{17}$$

That is, the pseudo-prior for $\boldsymbol{\theta}_j$ is given by the conditional posterior for $\boldsymbol{\theta}_j$ given observations of track $i = j$ and the pseudo-priors are independent for each j. Because of the

lower-triangular structure in Eq. (10) the pseudo-priors specified in Eq. (17) will always cover the conditional distribution for θ_j in the Gibbs update for the orbit selectors in Eq. (14).

The pseudo-prior for each orbit model is numerically defined by running separate MCMC chains for each orbit model to generate samples from $p(\boldsymbol{\theta}_j|\hat{\mathbf{y}}_j)$. This is a required pre-processing step before performing the sampling for model selection. Because the pseudo-prior for each orbit model is conditioned on the orbit selector parameters it is important to accurately estimate the pseudo-prior normalization from the MCMC samples, which typically requires many samples from a well-converged chain.

2.2.2. Complete sampling algorithm

Because the prior for $\boldsymbol{\Theta}$ depends on k, the full conditional distribution for k is no longer in a form allowing fast Gibbs samples to be drawn. Instead, we use Eq. (14) as a proposal distribution for MH updates of the orbit selectors. The modified sampling algorithm is:

1. Draw a proposal value of $k^{i'}$ for each i from the distribution in Eq. (14). Accept the proposal with MH acceptance probability,

$$\alpha_k = \frac{P(\boldsymbol{\Theta}|k', \hat{\mathbf{y}})}{P(\boldsymbol{\Theta}|k, \hat{\mathbf{y}})}. \tag{18}$$

This is just the ratio of $\boldsymbol{\Theta}$ priors because all other terms cancel due to the choice of proposal distribution.

2. For each i draw P_{orbit}^i as a Gibbs update from Eq. (15) (This step unchanged).

3. For each j propose a new $\boldsymbol{\theta}_j'$ using the statistical ranging proposal from Section 2.1 conditioned on times $t_i, t_{i'}$ with i, i' selected randomly from the set of integers $\left\{\ell | \ell = j, \ldots, N_{\mathrm{track}}; k_j^\ell = 1\right\}$. The proposal $\boldsymbol{\theta}_j'$ is accepted with probability,

$$\alpha_{\boldsymbol{\theta}} = \frac{L(\hat{\mathbf{y}}|\boldsymbol{\theta}_j')P(\boldsymbol{\theta}_j')}{L(\hat{\mathbf{y}}|\boldsymbol{\theta}_j)P(\boldsymbol{\theta}_j)} \frac{J(\boldsymbol{\theta}_j')}{J(\boldsymbol{\theta}_j)} \tag{19}$$

where $J(\boldsymbol{\theta}_j)$ is the Jacobian for the change of variables from the statistical ranging parameters to the orbital elements. When $k_j^i = 0 \,\forall i$ then $\boldsymbol{\theta}_j'$ is drawn from the pseudo-prior in Eq. (17) instead.

To improve the mixing of the chain, it can be helpful to iterate step 3 several times before moving to the next chain step.

Note that the sampling model developed to this point for linking orbit models is equivalent to the Reversible Jump MCMC (RJMCMC) algorithm for model selection (Green, 1995) when the proposal distribution for RJMCMC updates of the orbit parameters is chosen to be the pseudo-prior in Eq. (17). The equivalence of the two formalisms was first explained explicitly by Godsill (2001). To be efficient, the RJMCMC framework also requires a good choice

of proposal distribution for the orbit selectors, which is no
obvious in the current application. It is possible that usin
the RJMCMC formalism with a different proposal distr
bution could lead to even more efficient mixing than th
sampling algorithm described above. But we will not pu
sue this further because the algorithm above has so fa
been sufficient.

3. Case Study: LSST "visits"

The LSST observing cadence plan currently assume
that exposures will be 15 seconds. A single "visit" to
patch of the sky will consist of two exposures separate
by four seconds of time to read out the gigapixel CCI
Given that an object in GEO travels at an angular spee
as seen from the ground of about 15 arcseconds per second
GEO objects will be visible in LSST exposures as streak
about 225 arcseconds long. An LSST visit will yield tw
exposures with two closely spaced streaks giving four an
gular coordinates and times that can be used to find
POD. But, because 225 arcseconds is such a small fractio
of 360 degrees, the PODs of GEO objects will have larg
uncertainties. There is a plan for the LSST cadence t
revisit the same region of sky twice in one night with tin
separations of about 40-60 minutes.

So to understand how well LSST can determine GEO
orbits a key question to resolve is whether tracks from two
LSST visits separated by about 40 minutes can be reliably
linked together. This is the example we consider in this
section to demonstrate the POD and linking methods from
Section 2.

3.1. LSST visit Preliminary Orbit Determinations

First consider the POD from a single LSST visit. We
assume that two complete streaks are observed yielding RA
and DEC measurements at times $t = 0, 15, 19, 34$ seconds
from the start of the first exposure. We take a conserva-
tive estimate for the LSST seeing disk (0.8 arcseconds) to
be the 1-σ Gaussian error in the angular position measure-
ments so that Σ_i from Eq. (1) is

$$\Sigma_i = (0.8")^2 \mathbf{I}_4 \qquad (20)$$

We ran an MCMC chain with 5×10^5 steps using the like-
lihood of Eq. (1), statistical ranging proposals to update
the orbital elements, and stepping in the equinoctial pa-
rameters for the orbit. The trace plots of the chain steps
in Fig. 1 show that the statistical ranging proposals give
very fast mixing of the chain at some points, but can occa-
sionally get "stuck" at a single parameter value for many
consecutive steps. The regions where the chain is not mix-
ing well can be mitigated by running a very long chain,
or many independent chains, and then selecting a regular
subset of chain steps. By running two chains of 5×10^5
steps each and selecting every 400th step, We obtained
2500 uncorrelated samples from the posterior for the or-
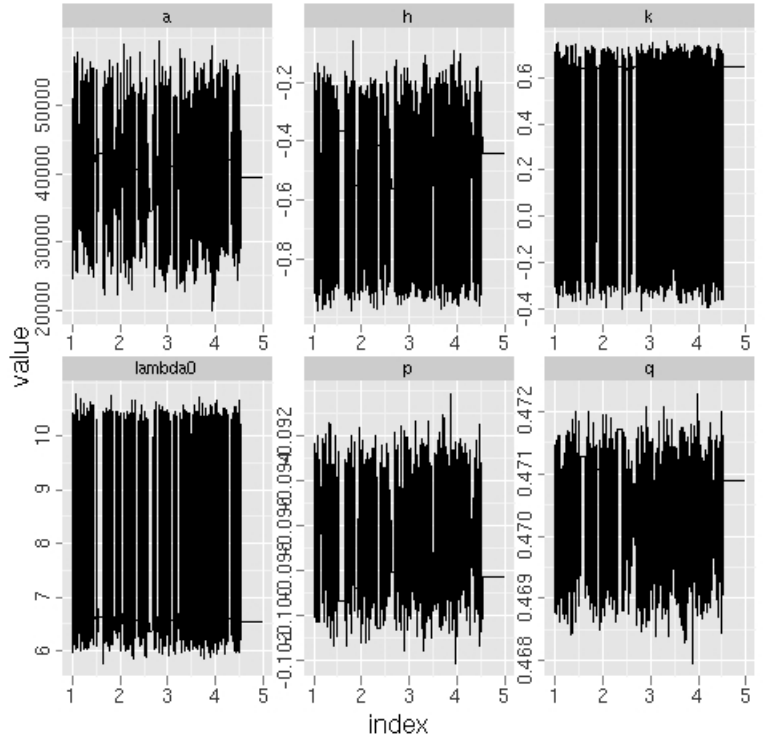bital elements given one LSST visit. In Fig. 2 we show

Figure 1: Trace plots of 500,000 MCMC samples for a single LSST
visit using statistical ranging proposals and MH updates.

scatterplots of all the two-dimensional projections of the
posterior samples. The vertical lines in each of the diag-
onal panels indicate orbital elements used to generate the
mock observations. The samples follow highly degenerate
and contorted distributions in the equinoctial parameter
space. First, this shows the advantage of using statistical
ranging proposals because most choices of proposal dis-
tributions in the equinoctial parameters would not sam-
ple the complicated posterior very efficiently. Second, the
maxima of the marginal posteriors for k and λ_0 are sig-
nificantly offset from the input model parameters. This is
a result of the projection of the degenerate multivariate
posterior distribution into one dimension and illustrates
how biased orbital elements could be inferred without a
full characterization of the joint posterior. Third, we can
see that the PODs from one LSST visit have large and, in
some cases, multimodal error distributions. It is clear that
the errors cannot be accurately described by a covariance
matrix (as in, e.g., Sabol et al., 2010).

To further demonstrate the benefits of the statistical
ranging proposal distribution in this case study, we show
the "path" taken by the MCMC chain in the h-k plane
in Figure 3, after thinning the chain by a factor of 100 to
remove highly correlated steps. The marginal posterior in
this plane follows a crescent shape, mostly concentrated in
two distinct maxima. A proposal distribution specified in
terms of the equinoctial elements would make it difficult
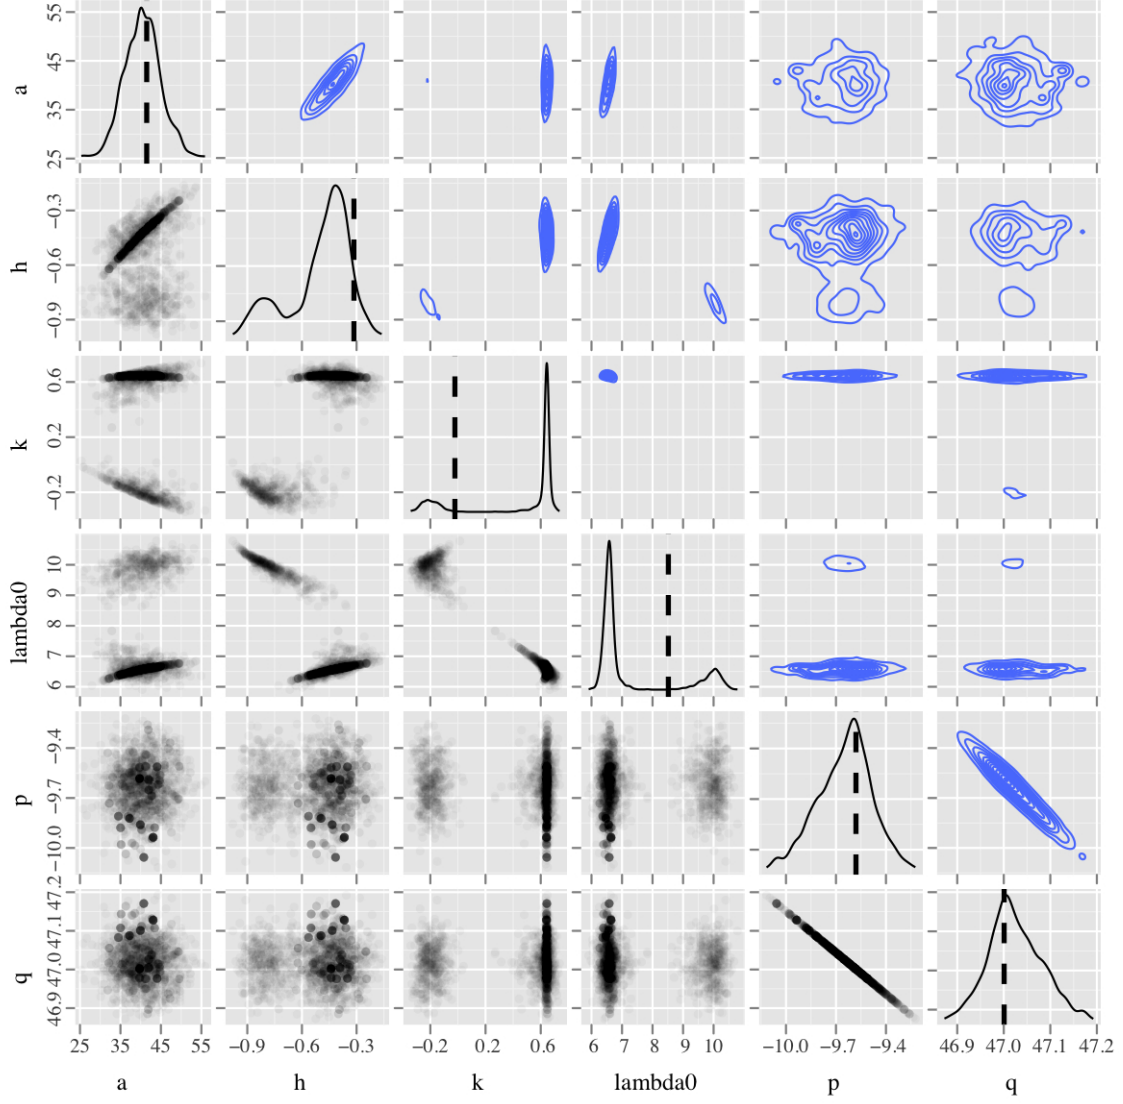to explore both maxima with a single chain, but the sta-

6

Figure 2: Posterior samples of the six equinoctial orbital elements given a single LSST visit. The diagonal panels show the marginal posterior probability distribution estimates for each of the equinoctial parameters. The vertical dashed lines in each diagonal panel indicate the input values for the mock data that was used to constrain the parameters. The lower-triangular panels show samples from the 2D marginal posteriors for parameter pairs where darker points indicate a larger density of samples. The upper-triangular panels are contour-plots of the same samples plotted in the lower-triangular panels.
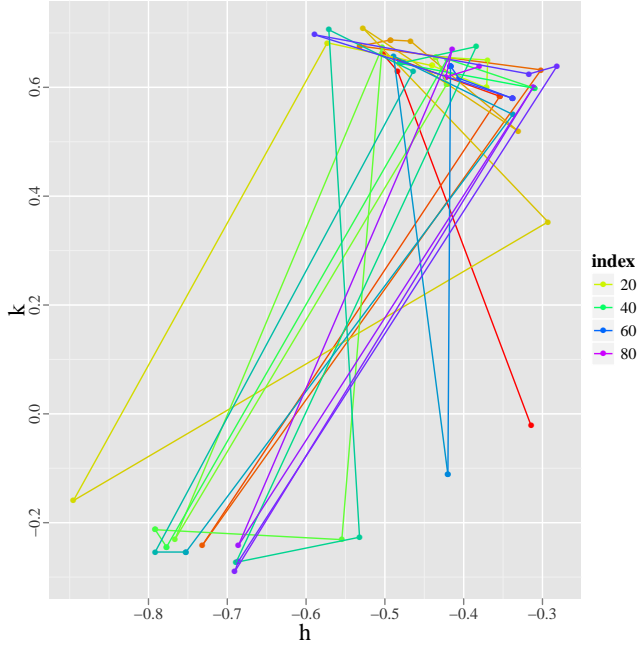
Figure 3: Sequential steps in the MCMC chain in the *h-k* parameter plane after thinning by a factor of 100 to remove correlations between chain steps. The colors indicate the chain step index increasing from yellow to purple.
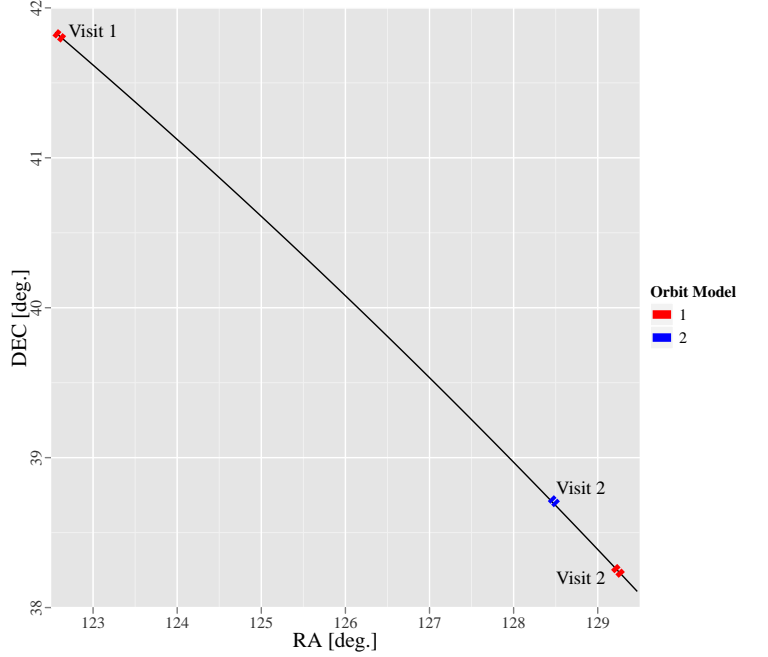

Figure 4: Angular positions of the mock track observations for two LSST visits. In the two scenarios considered, either orbit model 1 or 2 is observed in visit 2 (which is 40 minutes after visit 1). The solid black line shows the input orbit model 1. Because the blue tracks for orbit model 2 in visit 2 are so close to orbit model 1 in the sky, there is a significant chance of incorrectly linking orbit 1 to the blue tracks for orbit 2.

tistical ranging proposals move efficiently to sample the entire posterior. This ability to sample the orbital elements without significant burn-in and efficient mixing is crucial to the success of the model selection algorithm.

3.2. Linking tracks from two LSST visits

To test the algorithm for linking orbits, we consider a scenario with two LSST visits observed, separated by 40 minutes. For two visits there are either one or two distinct objects observed. The goals for the linking algorithm are to link the tracks in the two visits with high confidence when the tracks are generated by the same object and to reject the linking when the tracks are generated by distinct objects. The orbital parameters for the two distinct model objects are listed in Table 1.

The observed angular positions of the two orbit models are displayed in Fig. 4. The red lines show the observed streaks for orbit model 1 in both visits, while the blue lines show the streaks for orbit model 2 in the second visit only. Each visit has two separate lines denoting the streaks in the two 15 second exposures in each visit (which we assume are linked with absolute certainty). Even though orbit model 2 has very different semi-major axis and eccentricity from orbit model 1, the projected position on the sky in visit 2 is very close to the expected projected position for orbit model 1.

The marginal posterior probability distributions for both correct and incorrect linking are compared in Fig. 5. The red points and contours are copied from Fig. 2 and show the posterior samples when only the first observed track is

considered. The blue points and lines show the posterior when both the first and second visits are combined to refine the orbital elements of the same object seen in both visits. The green points and lines show the posterior samples for a *distinct* object constrained only by the track in the second observed visit. And, finally, the purple points and lines show the posterior distribution for the orbital elements constrained by both tracks when each track was generated by a distinct object. So, the purple lines demonstrate that for this example there is a non-negligible posterior probability to incorrectly link distinct objects observed in two LSST visits separated by 40 minutes. The true orbital elements for the two distinct objects in the scenario leading to the purple posterior distribution in Fig. 5 are the same as those yielding the red and green colored posteriors. In all projections in Fig. 5, the posterior for incorrectly linking distinct objects lies in the region of overlap of the single-visit posteriors for each object. So, if orbital element posteriors are estimated for each observed track in isolation, as suggested in Section 2.2.1 for defining pseudo-priors on the orbital elements, then it could be possible to quickly estimate which tracks have a possibility of being linked by evaluating the overlap regions of all single-track posterior combinations. This determination can be used to specify priors on the orbit-selector parameters. A final MCMC chain run in both the orbit-selector and orbital element parameters will determine the relative linking probabili-

Table 1: Orbital parameters for the two objects observed in two LSST visits.

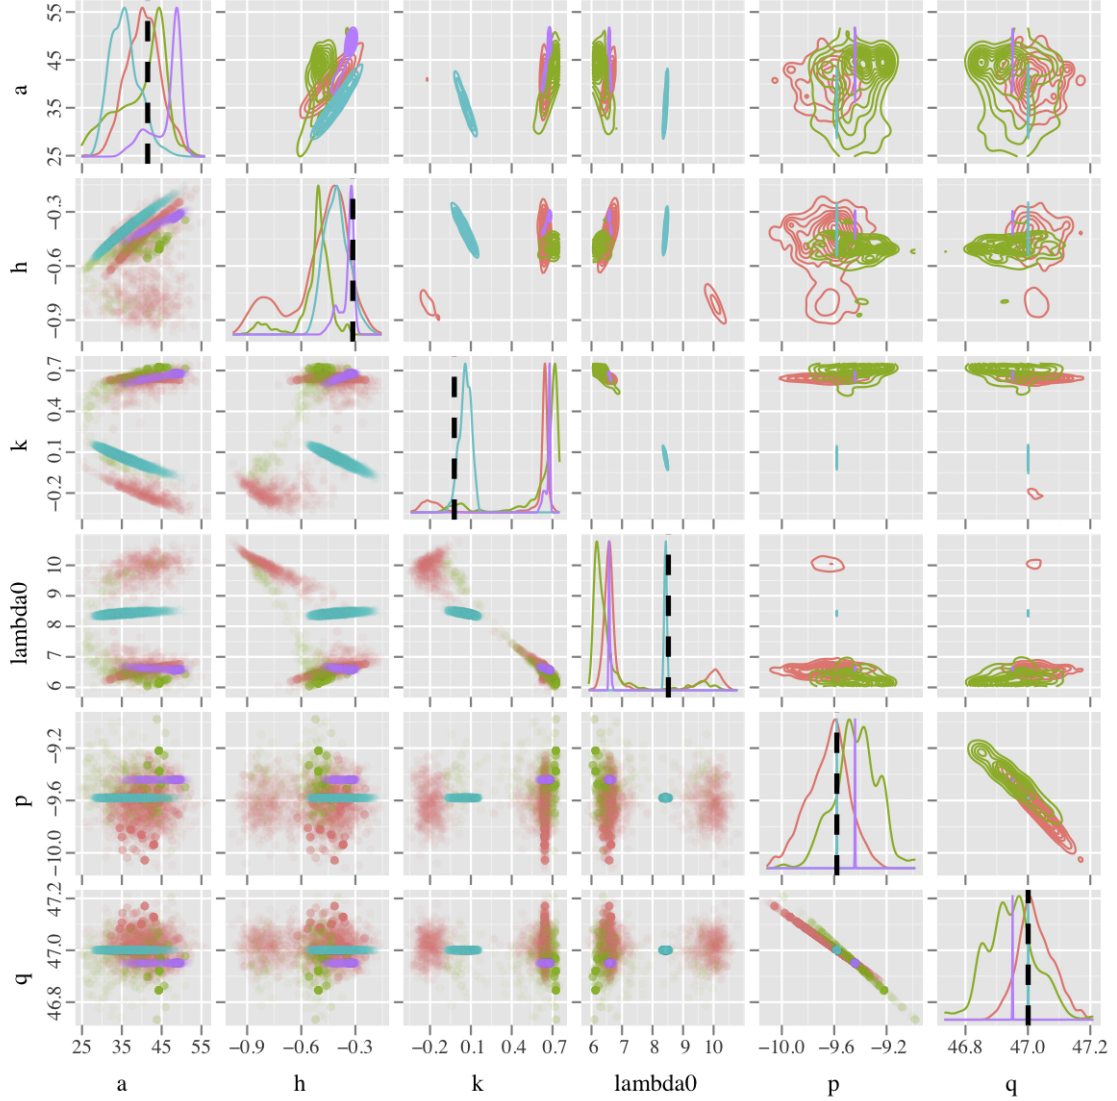| Model	semi-major axis (km)	eccentricity	period (days)	h	k	λ_0	p	q
1	41489	0.315	0.973	-0.314	-0.022	8.519	-0.096	0.470
2	42945	0.739	1.025	-0.369	0.640	6.649	-0.094	0.470

Figure 5: Two-dimensional projections of posterior samples of the orbital elements given observations from one or two LSST visits. The panels are arranged identically to Fig. 2 and the red points and lines are showing the same MCMC samples as in Fig. 2, which are samples from the posterior given only observations in the first LSST visit (object 1). The blue points and lines show the posterior given observations in two visits when the tracks in each visit correspond to the same object (object 1). (That is why the blue lines in the diagonal panels tighten around the vertical dashed lines denoting the "true" parameter values.) The green points and lines show the posterior given observations of a distinct object in the second of two visits (object 2). And the purple points and lines show the posterior given observations from two visits with the each visit observing a different object (objects 1 and 2). So, the purple points and lines are analogous to the blue points and lines but when the tracks in the two visits are erroneously linked.
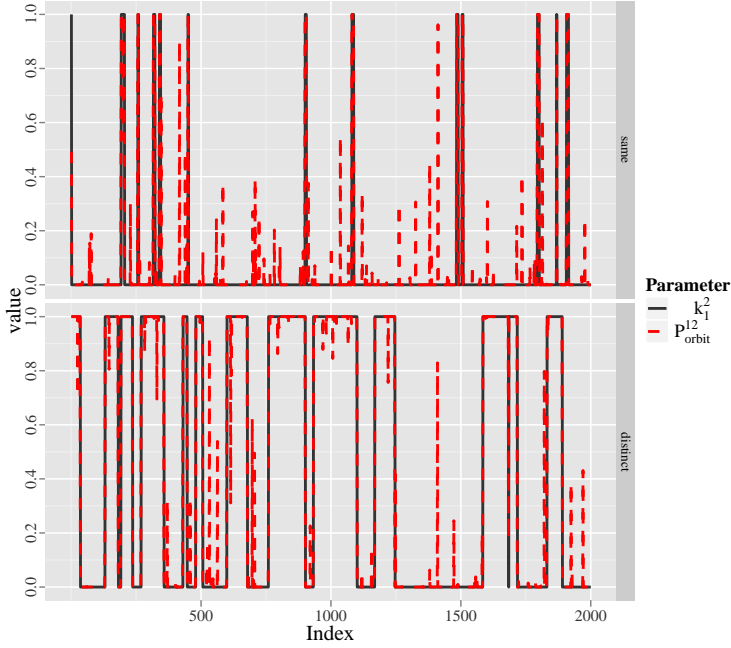
Figure 6: Trace plots of the orbit selector parameter k_1^2 and prior on the orbit selector P_{orbit}^{12} for the two scenarios when tracks are observed in two LSST visits that are generated by the "same" or "distinct" objects.

ties for each track combination without wasting computing time on tracks that have negligible probability to be linked.

Trace plots of the orbit selector samples for the two model scenarios are shown in Fig. 6. We show only the k_1^2 parameter because this uniquely specifies the assignment of the two tracks to the two orbit models. The P_{orbit}^{12} parameter samples are also shown by the red dashed lines. For the first scenario when both tracks are generated by the same orbit model, k_1^2 is predominantly zero with punctuated steps to a value of one showing that both possible track assignments were adequately considered by the chain. On the other hand, for the second scenario where each track was generated by a distinct orbit model, k_1^2 fluctuates more regularly between the values of zero and one. The P_{orbit}^{12} trace plots in each scenario show that many values of P_{orbit}^{12} are sampled during the chain, but k_1^2 only steps to a different value when P_{orbit}^{12} is very close to zero or one. The orbit selectors will step to new values more readily as the time interval between the two visits increases causing the orbit propagation uncertainty to increase accordingly.

The marginal posterior probability to link track 2 with orbit model 1 is

$$p(\mathrm{track2}|\mathrm{orbit1}) = \frac{1}{N_{\mathrm{MCMC}}} \sum_{i=1}^{N_{\mathrm{MCMC}}} k_{1,i}^2, \qquad (21)$$

where N_{MCMC} is the number of independent posterior samples from the MCMC chain. The results of evaluating

Table 2: Marginal posterior linking probabilities for the two generating orbit model scenarios.

Track generating model	$p(\mathrm{track2}|\mathrm{orbit1})$
Same (scenario 1)	0.97
Distinct (scenario 2)	0.45

Eq. (21) for the two scenarios in this case study are listed in Table 2. Our main result is that two distinct GEO objects observed in two LSST visits separated by 40 minutes can be erroneously linked with fairly large certainty (e.g., 45% for this case study) because a single LSST visit does not provide a sufficiently precise orbit determination.

4. Discussion and Conclusions

We have demonstrated a statistical sampling model to simultaneously constrain orbital elements and link uncorrelated tracks with a quantified linking probability. Our algorithm uses MCMC with statistical ranging proposals (Virtanen et al., 2001) to efficiently sample the joint posterior distribution for the orbital elements given a set of linked tracks in optical images. We also sample "orbit-selector" parameters defining which observed tracks are linked with which orbit models, where we assume that N_{track} tracks can be fit by up to N_{track} different orbit models. The posterior samples for the orbit-selector parameters from the MCMC chain determine the probability for the linking of all possible track combinations. Our algorithm is fundamentally different from many approaches that first perform a preliminary orbit determination (POD) and then propagate the orbit state vector to the times of subsequent observations. Instead we simultaneously constrain the parameters for the orbit given the full set of observations up to the current time. Viewed in this way, linking orbits is simply a parameter estimation problem, and not a problem in propagating the orbital state vector and its parameterized uncertainty.

A key ingredient in our sampling model is the definition of a "pseudo-prior" on the orbital elements for each possible orbit model that defines the MCMC updates when no tracks are assigned to a given orbit model. (The pseudo-prior is also functionally equivalent to the parameter proposal distribution in the Reversible Jump MCMC framework.) We find the orbit determination and linking works well when the pseudo-prior is the posterior probability for the orbital elements given a single observed track, as suggested by Godsill (2001). Pre-computing the pseudo-priors is equivalent to performing a POD for each observed track in turn. Because the statistical ranging proposals are so efficient, requiring essentially zero MCMC burn-in and yielding small correlations between chain steps, performing robust PODs is fast even for large numbers of objects. The KD-tree algorithms for linking uncorrelated tracks from Kubica et al. (2007); Granvik & Muinonen (2008) may be useful for speeding up the algorithm presented here if the

pre-computation of the pseudo-priors allows for a meaningful tree construction. We leave this investigation for future work.

Using the example of a single LSST "visit", defined to be two back-to-back 15 second exposures, we showed that the PODs for GEO objects can be multi-modal and highly degenerate in the six equinoctial parameters. When a distribution of GEO objects is observed, as expected with LSST, multiple linking solutions may be viable. In our case study of two LSST visits separated by 40 minutes, we found a probability of 45% to erroneously link two distinct objects. While this single example does not address how frequently tracks may be incorrectly linked with the LSST, it does demonstrate the probability to incorrectly link tracks can be quite large even with observations only 40 minutes apart. Our algorithm gives a probability that the linking is correct, allowing robust conclusions to be drawn about the GEO debris distribution given imperfect data. However, to constrain the debris distribution well enough for applications in, e.g. collision avoidance, either a dedicated campaign of real-time streak detection and follow-up or changes in the LSST observing cadence will likely be necessary.

In future work we plan to explore minor modifications to the LSST cadence that would allow useful GEO debris catalogues to be constructed without any follow-up observations. For example, if an LSST visit was defined to have more than two 15 second exposures then the PODs would improve and the probability for incorrect linking in subsequent visits would be reduced. Simply redefining a visit would destroy the observing cadence optimizations for other science goals. But it may be possible to use a small amount of time each night to repeatedly observe the same region of sky, with subsequent nights designed to fill out PODs for the entire GEO debris distribution. We also note that future work will need to incorporate non-conservative force models for propagating the orbits between nights. We can incorporate non-conservative forces by adding force model parameters to the list of orbit parameters to be constrained in exactly the same way as the equinoctial parameters are constrained in this paper.

Acknowledgments

Bowell, E., Virtanen, J., Muinonen, K., et al., *Asteroid Orbit Computation*, Asteroids III, 1, 27–43, 2002.

Broucke, R. A., Cefola, P. J., *On the Equinoctial Orbit Elements*, Celest. Mech., 5, 303–310, http://dx.doi.org/10.1007/BF01228432, 1972.

Carlin, B.P., Chib, S., *Bayesian model choice via Markov chain Monte Carlo methods*, J. Royal Stat. Soc. Series B (Methodological), 57, 3, 473–484, 1995.

DeMars, K., Jah, M., *Passive Multi-Target Tracking with Application to Orbit Determination for Geosynchronous Objects*, 19th AAS/AIAA Space Flight Mechanics Meeting, Savannah, Georgia, AAS Paper 09-108, 2009.

DeMars, K., Jah, M., Giza, D., et al., *Orbit Determination Performance for High Area-to-Mass Ratio Space Object Tracking Using an Adaptive Gaussian Mixtures Estimation Algorithm*, 21st International Symposium on Space Flight Dynamics, Toulouse, France, 2009.

DeMars, K., Jah, M., Schumacher, P., Jr., *The Use of Angle and Angle Rate Data for Deep-Space Orbit Determination and Track Association*, 20th AAS/AIAA Space Flight Mechanics Meeting, San Diego, CA, AAS Paper 10-153, 2010.

Fujimoto, K., Scheeres, D., *Correlation of Optical Observations of Earth-Orbiting Objects by Means of Probability Distributions*, paper presented at the 2010 AIAA/AAS Astrodynamics Specialist Conference, Toronot, August 2010. Paper AIAA-2010-7975. Accepted for publication in the Journal of Guidance, Control and Dynamics.

Godsill, S.J., *On the relationship between Markov chain Monte Carlo methods for model uncertainty*, J. Comp. Graph. Stat., 2, 230–248, 10, 2001.

Granvik, M., *Asteroid identification using statistical orbital inversion methods*, Doctoral dissertation, University of Helsinki, Faculty of Science, Department of Astronomy, http://urn.fi/URN:ISBN:978-952-10-4409-0, 2007.

Granvik, M., Muinonen, K., *Asteroid identification over apparitions*, Icarus, 198, 1, 130–137, DOI: 10.1016/j.icarus.2008.06.005, 2008.

Green, P.J., *Reversible jump Markov chain Monte Carlo computation and Bayesian model determination*, Biometrika, 4, 711, 82, 1995.

Gronchi, G. F., *Classical and modern orbit determination for asteroids*, Proc. Int. Astron. Union, 293-303, 2004.

Hogg, D. W., Bovy, J., Lang, D., *Data analysis recipes: Fitting a model to data*, ArXiv e-prints, 2010.

Kubica, J., Denneau, L., Grav, T., et al., *Efficient intra- and inter-night linking of asteroid detections using kd-trees*, Icarus, 189, 151–168, http://dx.doi.org/10.1016/j.icarus.2007.01.008, 2007.

LSST Science Collaborations and LSST Project, *LSST Science Book*, 2nd Edition, ArXiv e-prints, 0912.0201, http://www.lsst.org/lsst/scibook, 2009.

Marsden, B. G., *The computation of orbits in indeterminate and uncertain cases*, Astron. J., 102, 1539-1552, http://dx.doi.org/10.1086/115979, 1991.

Maruskin, J., Scheeres, D., Alfriend, K., *Correlation of Optical Observations of Objects in Earth Orbit*, J. Guid. Contr. Dyn., 0731-5090, 32, 1, 194–209, doi: 10.2514/1.36398 2009.

Milani, A., Gronchi, G., Vitturi, M., et al., *Orbit determination with very short arcs. I admissible regions* Cel. Mech. Dyn. Astron., 90, 1, 57–85, DOI: 10.1007/s10569-004-6593-5, 2004.

Muinonen, K., Bowell, E., *Asteroid Orbit Determination Using Bayesian Probabilities*, Icarus, 2, 104, 255 - 279, http://dx.doi.org/10.1006/icar.1993.1100, 1993.

Oszkiewicz, D., Muinonen, K., Virtanen, J. et al., *Asteroid orbital ranging using Markov-Chain Monte Carlo*, Meteoritics & Planetary Science, 12, 44, 1897–1904, http://dx.doi.org/10.1111/j.1945-5100.2009.tb01999.x, 2009.

Payne, T., Gregory, S. A., Tombasco, J., et al., *Satellite Monitoring, Change Detection, and Characterization Using Non-Resolved Electro-Optical Data from a Small Aperture Telescope*, Advanced Maui Optical and Space Surveillance Technologies Conference, 2007.

Sabol, C., Sukut, T., Alfriend, K. T. et al., *Linearized Orbit Covariance Generation and Propagation Analysis via Simple Monte*

Carlo Simulations, 20th AAS/AIAA Space Flight Mechanics Meeting, San Diego, CA, AAS Paper 10-134, 2010.

Schildknecht, T., Ploner, M., Hugentobler, U., *The search for debris in GEO*, Adv. Space Res., 9, 28, 1291 - 1299, http://dx.doi.org/10.1016/S0273-1177(01)00399-4, 2001.

Schildknecht, T., Musci, R., Ploner, M. et al., *Optical observations of space debris in GEO and in highly-eccentric orbits*, Adv. Space Res., 5, 34, 901 - 911, http://dx.doi.org/10.1016/j.asr.2003.01.009, 2004.

Shefer, VA, *New method of orbit determination from two position vectors based on solving Gauss's equations*, Sol. Sys. Res., 3, 44, 252–266, 2010.

Tommei, G., Milani, A., Rossi, A., *Orbit determination of space debris: admissible regions*, Celest. Mech. Dyn. Astron., 4, 97, 289–304, DOI: 10.1007/s10569-007-9065-x, 2007.

Virtanen, J., Muinonen, K., Bowell, E., *Statistical ranging of asteroid orbits*, Icarus, 2, 154, 412–431, 2001.